NCVHS Testimony – Epic Systems Corporation

Introduction

Recognized as an innovative technology company with client-centered values, Epic Systems Corporation is a leader in developing and supporting state-of-the-art, feature-rich software that helps organizations change the way healthcare is delivered. Its comprehensive portfolio of clinical, access, and financial systems centers around a single patient record, delivering data-level integration across *all* care environments and providing the information management that organizations need to enhance services, trim costs, and delight patients. Every day, Epic's systems help many of the country's leading healthcare organizations take care of patients — and their business.

Based in Madison, Wisconsin, Epic was founded in 1979 by a group of healthcare and IT professionals. Today, Epic remains one of the largest and most successful privately owned vendors of healthcare information systems.

Answers to Testimony Questions

1) Do the PMRI message format standards on the attached list address your highest priorities for exchanging PMRI electronically?

   Yes, for most practical purposes.

   a) If not, what priorities of yours are not met by these standards?

      We are concerned with the lack of an overall PMRI structure among the discussed standards. It is our belief that the listed standards provide the means for exchanging *pieces* of the PMRI structure, but that they do not provide consistent overall framework. We believe that HL7 v3 provides such a global framework.

2) What is the name of the specific standard and version that your comments will address in response to questions 3 through 6? (If you wish to make comments on more than one specific standard or version please do so with a separate contiguous listing of questions 2 through 6 so that we can relate each of your responses to the appropriate standard or version.)

   The following standards will be addressed below:

   - HL7 Ver. 2.x Order Entry Transaction Set

   - HL7 Ver. 2.x Observation Reporting Transaction Set

   - HL7 Ver. 2.x Patient Administration Financial Management Transaction Set

   - HL7 Ver. 2.x Patient Administration Transaction Set

   - HL7 Ver. 2.x Scheduling Transaction Set

   - HL7 Ver. 2.x Medical Record/Image Management Transaction Set

   - HL7 Ver. 3 Clinical Document Architecture Framework

HL7 Ver. 2.x Order Entry Transaction Set and
HL7 Ver. 2.x Observation Reporting Transaction Set

3) What functions or characteristics caused you to select this standard for implementation?

(Chapters 4 and 7, respectively)

Providing a rich order entry and results review functionality has been crucial for the success of our clinical applications since their introduction in 1994. The two transaction sets discussed here have allowed us to provide this functionality to customers such as Kaiser Permanente NW, Harvard Vanguard Medical Associates, and other large groups. These sets have facilitated a successful exchange of clinical information between our applications and other external entities.

These particular standards were not as popular or as accepted in 1994 as they are now. However, we saw their potential and understood that direct participation on our part was crucial to developing a promising standard.

We selected these specific transaction sets based on:

- Relative completeness of data elements included in the standard.

- Coverage of representative workflows for our customers.

- Growing acceptance of the transaction sets among vendors.

- Data representation in an abstract, application-neutral format.

4) What was your experience implementing this standard?

Our answers to the following questions cover our view and efforts in developing the general framework needed to implement the HL7 standard as part of our application functionality. Our answers do not, however, reflect the possible effort that our customers needed to utilize the standard for actual data exchange among vendors.

Along with incorporating data elements included in the standard within our clinical applications, we implemented the general framework (hierarchical model) of HL7 messaging. Based on this generic framework, implementation of the various HL7 transaction sets was relatively simple.

a) What aspect of implementing this standard exceeded your expectations?

None.

b) How long did it take for your organization to implement this standard?

The initial interface development following this standard took us approximately six months (both interfaces). However, interface development is an ongoing process as the standard evolves and as we add more options/functionality to our interfaces and applications.

When we implement these interfaces with our customers today, the usual implementation time (analysis, installation of the interface, application setup, custom programming, and testing) is between three and six months, depending on the complexity of the overall project.

c) What was the total implementation cost? (Please include your best estimate for fees, education, interface engines, custom coding, testing, maintenance, etc.)

The implementation cost for customers consists of license fees, staffing for the implementation period, and maintenance staffing once the interface is active in a real data environment. These costs vary based on the complexity of a customer's system and specific interface needs.

d) After you completed the implementation of this standard, did it prove to be easy to use?

The main reasons for the relatively long implementation and testing time are:

- Not all vendors follow the standard and very often we have to compensate for that.

- Ambiguity and optionality in the standard's specifications.

It is our belief that the main reason for the ambiguity and optionality in the specifications comes from the fact that the individual transactions we address here were not designed to be part of a global PMR framework.

For example, the current version of the standard, "suggests" a "preferred" way to transmit microbiology results (HL7 version 2.4, Chapter 7, p. 59). The suggested format is very precise and contains all of the required information for a complete microbiology result transmission. However, the fact that this part of the specifications was not made required has resulted in many vendors implementing microbiology results reporting in an ad-hoc and less precise way.

5) What are the major weaknesses or limitations of this standard(s)?

    a) What could or should be done to improve or strengthen this standard?

The standard should be based on a solid foundation using advanced methods and tools in modeling and systems design. This will ensure internal consistency among transaction sets and will greatly improve interoperability. HL7 has realized this requirement and has developed a Reference Information Model (RIM) as the basis of its new version 3 efforts. Currently the version 3 of the HL7 transaction sets is under ballot, and we believe this new version should be considered by the committee for inclusion in any future recommendations.

6) Have you evaluated any more current version(s) of the standards you are using?

    a) If so, what is your evaluation of them?

In the past three years, we have participated in the HL7 demos at the annual HIMSS expositions. We have developed prototypes for version 3 of these transaction sets, each time adapting our prototypes to the then current format and specifications. We've found it relatively easy to implement because of the following features of the new standard:

- The format is based on the Extensible Markup Language (XML), and we can use freely-available tools in our prototype implementation.

- The underlying Reference Information Model (RIM) facilitates the reuse of computational objects (source code, programming libraries, configuration settings, etc.), and provides a consistent framework for implementing these interfaces.

We are currently actively participating in the balloting of the new version 3, and we plan to implement the new version as soon as it is approved (the expected timeframe for completing the balloting process is the middle of 2002).

HL7 Ver. 2.x Patient Administration Financial Management Transaction Set

3) What functions or characteristics caused you to select this standard for implementation?

(Chapter 6)

This transaction set provides an acceptable standard for the exchange of clinically related financial data.

The popularity and acceptance of this particular standard was relatively high when we adopted it (around 1993) and provided us with a way to exchange this type of data with a number of other systems, which was and has been a benefit for many of our customers.

We selected this specific transaction set based on:

- Relative completeness of data elements included in the standard.

- Coverage of representative workflows for our customers.

- Growing acceptance of the transaction sets among vendors.

- Data representation in an abstract, application-neutral format.

4) What was your experience implementing this standard?

Our answers to the following questions cover our view and efforts in developing the general framework needed to implement the HL7 standard as part of our application functionality. Our answers do not, however, reflect the possible effort that our customers needed to utilize the standard for actual data exchange among vendors.

Along with incorporating data elements included in the standard in our clinical applications, we implemented the general framework (hierarchical model) of HL7 messaging. Based on this generic framework, implementation of the various HL7 transaction sets was relatively simple.

a) What aspect of implementing this standard exceeded your expectations?

None.

b) How long did it take for your organization to implement this standard?

The initial interface development following this standard took us approximately six months. However, the interface development is an ongoing process as the standard evolves and as we add more options/functionality to our interfaces and applications.

When we implement these interfaces with our customers, the usual implementation time (analysis, installation of the interface, application setup, custom programming, testing) is between three and six months, depending on the complexity of the overall project.

c) What was the total implementation cost? (Please include your best estimate for fees, education, interface engines, custom coding, testing, maintenance, etc.)

The implementation cost for customers consists of license fees, staffing for the implementation period, and maintenance staffing once the interface is active in a real data environment. These costs vary based on the complexity of a customer's system and specific interface needs.

d) After you completed the implementation of this standard, did it prove to be easy to use?

The main reasons for the relatively long implementation and testing time are:

- Not all vendors follow the standard and very often we have to compensate for that.
- Ambiguity and optionality in the standard's specifications.

  It is our belief that the main reason for the ambiguity and optionality in the specifications comes from the fact that the individual transactions we address here were not designed to be part of a global PMR framework.

  For example, the DFT^P03 message type (chapter 6) contains a DG1 segment, a PR1 segment, and an FT1 segment. The purpose of the DG1 and PR1 segments is to specify the diagnoses and procedures related to the financial transaction in the FT1 segment. At the same time, the FT1 segment itself contains fields for both the diagnoses and the procedures associated with the financial transaction.

5) What are the major weaknesses or limitations of this standard(s)?

a) What could or should be done to improve or strengthen this standard?

The standard should be based on a solid foundation using advanced methods and tools in modeling and systems design. This will ensure internal consistency among transaction sets and will greatly improve interoperability. HL7 has realized this requirement and has developed a Reference Information Model (RIM) as the basis of its new version 3 efforts. Currently the version 3 of the HL7 transaction sets is under ballot, and we believe this new version should be considered by the committee for inclusion in any future recommendations.

6) Have you evaluated any more current version(s) of the standards you are using?

a) If so, what is your evaluation of them?

In the version 3 of the standard, the RIM provides a rich and consistent foundation for exchanging the type of data contained in this transaction set.

Currently, the ballot for version 3 is not yet completed. It addresses the data relationship between patients, clinical, and financial data better than the previous versions of the standard.

We are actively participating in the balloting of the new version 3, and we plan to implement the new version as soon as it is approved (the expected timeframe for completing the balloting process is the middle of 2002).

HL7 Ver. 2.x Patient Administration Transaction Set

3) What functions or characteristics caused you to select this standard for implementation?

(Chapter 3)

This transaction set provides an acceptable standard for the exchange of patient demographic and visit-specific data.

The popularity and acceptance of this particular standard was relatively high when we adopted it (around 1993) and provided us with a way to exchange this type of data with a number of other systems, which was and has been a benefit for many of our customers.

We selected this specific transaction set based on:

- Relative completeness of data elements included in the standard.

- Coverage of representative workflows for our customers.

- Growing acceptance of the transaction sets among vendors.

- Data representation in an abstract, application-neutral format.

4) What was your experience implementing this standard?

Our answers to the following questions cover our view and efforts in developing the general framework needed to implement the HL7 standard as part of our application functionality. Our answers do not, however, reflect the possible effort that our customers needed to utilize the standard for actual data exchange among vendors.

Along with incorporating data elements included in the standard in our clinical applications, we implemented the general framework (hierarchical model) of HL7 messaging. Based on this generic framework, implementation of the various HL7 transaction sets was relatively simple.

a) What aspect of implementing this standard exceeded your expectations?

None.

b) How long did it take for your organization to implement this standard?

The initial interface development following this standard took us approximately six months. However, the interface development is an ongoing process as the standard evolves and as we add more options/functionality to our interfaces and applications.

When we implement these interfaces with our customers, the usual implementation time (analysis, installation of the interface, application setup, custom programming, testing) is between two and four months, depending on the complexity of the overall project.

c) What was the total implementation cost? (Please include your best estimate for fees, education, interface engines, custom coding, testing, maintenance, etc.)

The implementation cost for customers consists of license fees, staffing for the implementation period, and maintenance staffing once the interface is active in a real data environment. These costs vary based on the complexity of a customer's system and specific interface needs.

d) After you completed the implementation of this standard, did it prove to be easy to use?

The main reasons for the relatively long implementation and testing time are:

- Not all vendors follow the standard and very often we have to compensate for that.

- Ambiguity and optionality in the standard's specifications.

  It is our belief that the main reason for the ambiguity and optionality in the specifications comes from the fact that the individual transactions we address here were not designed to be part of a global PMR framework.

  The lack of a place to specify the patient's Primary Care Provider (PCP) before version 2.3 is another example of the lack of a comprehensive approach to the PMRI concept.

5) What are the major weaknesses or limitations of this standard(s)

   a) What could or should be done to improve or strengthen this standard?

   The standard should be based on a solid foundation using advanced methods and tools in modeling and systems design. This will ensure internal consistency among transaction sets and will greatly improve interoperability. HL7 has realized this requirement and has developed a Reference Information Model (RIM) as the basis of its new version 3 efforts. Currently the version 3 of the HL7 transaction sets is under ballot, and we believe this new version should be considered by the committee for inclusion in any future recommendations.

6) Have you evaluated any more current version(s) of the standards you are using?

   a) If so, what is your evaluation of them?

   In the past three years, we have participated in the HL7 demos at the annual HIMSS expositions. We have developed prototypes for version 3 of these transaction sets, each time adapting our prototypes to the then current format and specifications. We've found it relatively easy to implement, because of the following features of the new standard:

   - The format is based on the Extensible Markup Language (XML), and we can use freely available tools in our prototype implementation.

   - The underlying Reference Information Model (RIM) facilitates the reuse of computational objects (source code, programming libraries, configuration settings, etc.), and provides a consistent framework for implementing these interfaces.

   We are currently actively participating in the balloting of the new version 3, and we plan to implement the new version as soon as it is approved (the expected timeframe for completing the balloting process is the middle of 2002).

HL7 Ver. 2.x Scheduling Transaction Set

3) What functions or characteristics caused you to select this standard for implementation?

   (Chapter 10)

   This transaction set provides a fairly rich standard for the exchange of appointment data.

   We were one of the first vendors to implement this transaction set for the electronic exchange of appointment data (when this chapter was added in version 2.3).

   We selected this specific transaction set based on:

   - Relative completeness of data elements included in the standard.

   - Coverage of representative workflows for our customers.

   - Growing acceptance of the transaction sets among vendors.

   - Data representation in an abstract, application-neutral format.

   - Similarities in the structure of the messages with already existing transaction sets, which allowed us to reuse some of our existing work.

4) What was your experience implementing this standard?

Our answers to the following questions cover our view and efforts in developing the general framework needed to implement the HL7 standard as part of our application functionality. Our answers do not, however, reflect the possible effort that our customers needed to utilize the standard for actual data exchange among vendors.

Along with incorporating data elements included in the standard in our clinical applications, we implemented the general framework (hierarchical model) of HL7 messaging. Based on this generic framework, implementation of the various HL7 transaction sets was relatively simple.

a) What aspect of implementing this standard exceeded your expectations?

None.

b) How long did it take for your organization to implement this standard?

The initial interface development following this standard took us approximately six months. However, the interface development is an ongoing process as the standard evolves and as we add more options/functionality to our interfaces and applications.

When we implement these interfaces with our customers, the usual implementation time (analysis, installation of the interface, application setup, custom programming, testing) is between two and four months, depending on the complexity of the overall project.

c) What was the total implementation cost? (Please include your best estimate for fees, education, interface engines, custom coding, testing, maintenance, etc.)

The implementation cost for customers consists of license fees, staffing for the implementation period, and maintenance staffing once the interface is active in a real data environment. These costs vary based on the complexity of a customer's system and specific interface needs.

d) After you completed the implementation of this standard, did it prove to be easy to use?

The main reasons for the relatively long implementation and testing time are:

- Not all vendors follow the standard and very often we have to compensate for that.
- Ambiguity and optionality in the standard's specifications.

  It is our belief that the main reason for the ambiguity and optionality in the specifications comes from the fact that the individual transactions we address here were not designed to be part of a global PMR framework.

5) What are the major weaknesses or limitations of this standard(s)?

a) What could or should be done to improve or strengthen this standard?

The standard should be based on a solid foundation using advanced methods and tools in modeling and systems design. This will ensure internal consistency among transaction sets and will greatly improve interoperability. HL7 has realized this requirement and has developed a Reference Information Model (RIM) as the basis of its new version 3 efforts. Currently the version 3 of the HL7 transaction sets is under ballot, and we believe this new version should be considered by the committee for inclusion in any future recommendations.

6) Have you evaluated any more current version(s) of the standards you are using?

a) If so, what is your evaluation of them?

In the version 3 of the standard, the RIM provides a rich and consistent foundation for exchanging the type of data contained in this transaction set.

Currently, the ballot for version 3 is not yet completed. It addresses the data relationship between patients, clinical, and financial data better than the previous versions of the standard.

We are actively participating in the balloting of the new version 3, and we plan to implement the new version as soon as it is approved (the expected timeframe for completing the balloting process is the middle of 2002).

HL7 Ver. 2.x Medical Record/Image Management Transaction Set

3) What functions or characteristics caused you to select this standard for implementation?

(Chapter 9)

This transaction set provides a fairly rich standard for the exchange of medical document management information.

We were one of the first vendors to implement this transaction set for the electronic exchange of medical transcription data (when this chapter was added in version 2.3).

We selected this specific transaction set based on:

- Relative completeness of data elements included in the standard.

- Coverage of representative workflows for our customers.

- Expected acceptance of the transaction sets among vendors.

- Data representation in an abstract, application-neutral format.

- Similarities in the structure of the messages with already existing transaction sets, which allowed us to reuse some of our existing work.

4) What was your experience implementing this standard?

Our answers to the following questions cover our view and efforts in developing the general framework needed to implement the HL7 standard as part of our application functionality. Our answers do not, however, reflect the possible effort that our customers needed to utilize the standard for actual data exchange among vendors.

Along with incorporating data elements included in the standard in our clinical applications, we implemented the general framework (hierarchical model) of HL7 messaging. Based on this generic framework, implementation of the various HL7 transaction sets was relatively simple.

a) What aspect of implementing this standard exceeded your expectations?

None.

b) How long did it take for your organization to implement this standard?

The initial interface development following this standard took us approximately six months. However, the interface development is an ongoing process as the standard evolves and as we add more options/functionality to our interfaces and applications.

When we implement these interfaces with our customers, the usual implementation time (analysis, installation of the interface, application setup, custom programming, testing) is between two and four months, depending on the complexity of the overall project.

c) What was the total implementation cost? (Please include your best estimate for fees, education, interface engines, custom coding, testing, maintenance, etc.)

The implementation cost for customers consists of license fees, staffing for the implementation period, and maintenance staffing once the interface is active in a real data environment. These costs vary based on the complexity of a customer's system and specific interface needs.

d) After you completed the implementation of this standard, did it prove to be easy to use?

The main reasons for the relatively long implementation and testing time are:

- Not all vendors follow the standard and very often we have to compensate for that. Furthermore, currently this specific set of transactions is not as widely used as the rest of the transaction sets discussed in this document.

- Ambiguity and optionality in the standard's specifications.

It is our belief that the main reason for the ambiguity and optionality in the specifications comes from the fact that the individual transactions we address here were not designed to be part of a global PMR framework.

According to this transaction set, a minimum of two messages are necessary for the creation of a new order (initiated from an external system) with an arbitrary test and the filing of a transcription to this new order (an ORM message followed by an MDM message). Only relatively recently has a new proposal been discussed which allows the above functionality within one message.

5) What are the major weaknesses or limitations of this standard(s)?

a) What could or should be done to improve or strengthen this standard?

The standard should be based on a solid foundation using advanced methods and tools in modeling and systems design. This will ensure internal consistency among transaction sets and will greatly improve interoperability. HL7 has realized this requirement and has developed a Reference Information Model (RIM) as the basis of its new version 3 efforts. Currently the version 3 of the HL7 transaction sets is under ballot, and we believe this new version should be considered by the committee for inclusion in any future recommendations.

6) Have you evaluated any more current version(s) of the standards you are using?

a) If so, what is your evaluation of them?

In the version 3 of the standard, the RIM provides a rich and consistent foundation for exchanging the type of data contained in this transaction set.

Currently, the ballot for version 3 is not yet completed. It addresses the data relationship between patients, clinical, and financial data better than the previous versions of the standard.

We are actively participating in the balloting of the new version 3, and we plan to implement the new version as soon as it is approved (the expected timeframe for completing the balloting process is the middle of 2002).

HL7 Ver. 3 Clinical Document Architecture Framework

3) What functions or characteristics caused you to select this standard for implementation?

The HL7 Clinical Document Architecture defines markup for clinical documents, based on the artifacts that exist in today's healthcare environment. The standard provides an important uniformity for the disparate document formats existing today, while allowing local expressiveness to be preserved.

We followed, and continue to follow, the development of the standard, and implemented a prototype for the HL7 HIMSS 2001 demo. As a vendor of a comprehensive computerized patient medical record system, we offer varied clinical documentation functionality in our applications. This particular standard can complement this functionality by allowing our customers to express this documentation in a common format.

4) What was your experience implementing this standard?

Our answers to the following questions cover our view and efforts in developing a prototype. We estimate that a full implementation of the standard will require more time and effort.

a) What aspect of implementing this standard exceeded your expectations?

Since this standard is part of the HL7 version 3 effort, and it is built based on the common HL7 Reference Information Model, we found that we could very easily reuse the parts already developed for the HL7 version 3 messages, which were also part of the prototype demo implementation.

b) How long did it take for your organization to implement this standard?

The prototype development took us approximately one month. As already stated, we expect a full implementation to take significantly more time and effort, as this standard is much more closely related to some of the core functionality of our applications.

c) What was the total implementation cost? (Please include your best estimate for fees, education, interface engines, custom coding, testing, maintenance, etc.)

N/A

d) After you completed the implementation of this standard, did it prove to be easy to use?

During the HL7 demo, documents created using this standard were transmitted to several different applications, and they were able to properly process and display the contents. Even though this was done in a demo setting, the ease of use of the generated documents seemed to exceed that of using the existing version 2.x prototypes.

5) What are the major weaknesses or limitations of this standard(s)

a) What could or should be done to improve or strengthen this standard?

This standard is part of the HL7 version 3 suite of standards. As the rest of these standards become normative, the Clinical Documentation Architecture Framework will continue to evolve and benefit from the work done in the HL7 Reference Information Model. An important undertaking of HL7 is the development of a document ontology as related to healthcare, which is part of the LOINC code set. Completing this task, and incorporating it in the CDA will strengthen the standard and increase interoperability.

6) Have you evaluated any more current version(s) of the standards you are using?

a) If so, what is your evaluation of them?

This is the latest version of the standard. We have developed a prototype as part of our participation in the HL7 demo at HIMSS 2001. We found it relatively easy to implement, because of the following features of this standard:

- The format is based on the Extensible Markup Language (XML), and we can use freely available tools in our prototype implementation.

- The underlying Reference Information Model (RIM) facilitates the reuse of computational objects (source code, programming libraries, configuration settings, etc.), and provides a consistent framework for implementing these interfaces.

We are currently actively participating in the further development of the standard, and we are preparing to start full implementation as soon as the rest of version 3 is approved (the expected timeframe for completing the balloting process is the middle of 2002).

7) Are there any other standards that support exchange of patient medical record information that NCVHS should be considering?

As it should be fairly obvious from our testimony, we believe that the main standard for exchange of patient medical record information should be HL7 version 3. There are several important reasons behind this:

- Version 3 of the standard is based on a comprehensive abstract reference information model. The development of this model has been done over a significant period of time, and with the participation of almost every major player in the healthcare information technology field: all major healthcare information systems vendors, leading healthcare providers like Kaiser Permanente and the Mayo Clinic, consulting companies like Cap Gemini Ernst and Young and the Gartner Group, government agencies like the VA and the Department of Defense. The methodology employed in the development in the version 3 reference information model (RIM) will bring healthcare information technology into the 21st century. Because of the comprehensive nature of the model, issues like security can be addressed a part of the standard, rather than being treated as add-ons at a later point. HL7 is also an international organization with affiliates around the world. The RIM is the result of taking into account and ensuring compatibility with existing work in other parts of the world like the CEN effort in Europe and the Government Electronic Health Record in Australia.

- The standard's methodology requires formal descriptions of application roles and interactions for every transaction set. This provides a certain level of conformance requirements, which is not available in previous versions. When the HL7 version 2.x transactions were made part of the Claims Attachments provisions of HIPAA, there was a significant effort needed to strictly define the constraints under which the HL7 transactions are deemed as satisfying HIPAA's requirements.

- The format used by HL7 version 3 is XML. This allows implementers of the standard to use freely existing tools to process the different transactions. This results in increased interoperability, quicker adoption, and lower cost of implementation. In addition, HL7 version 3 provides another very important feature: format independence. Because of the Reference Information Model underlying the standard, the semantics represented by the transaction sets are independent from the format. While currently XML provides the best format for these transactions, if a better one is invented, it can be easily adopted by the standard without changing the meaning already expressed in existing transaction sets.

There may be several possible arguments against considering HL7 version 3 for exchange of patient medical record information. We would like to address the following:

- *The development of the new version has not been completed yet*. We would like to point out that the development of this version started in 1996, and the current reference information model (RIM), as well as the development methodology, are rather mature. The new version is under ballot now, and it is expected to be approved by the middle of 2002. Our understanding of the scope and goals of adopting standards for the exchange of patient medical record information is that the new version will be available well in advance of the adoption of any specific requirements.

- *Version 3 has no installed base, while version 2.x is widely adopted*. It is very educational to look at the process of including the HL7 version 2.3.1 transaction sets as part of the requirements for Claims Attachments under HIPAA. Even though the transaction sets had existed for several years and were widely adopted, the specificity needed for that particular purpose required significant effort in removing ambiguity and optionality from the standard. The result has been a transaction set that requires just as much effort in implementation as a totally new transaction set. In contrast, version 3 of the standard incorporates features within its development methodology that will enable a specification close to, if not identical to, the ones used by organizations in cases falling outside of the realm of exchanging patient medical record information.

- *Since version 3 uses a different format and methodology from version 2.x, making it a requirement will significantly increase implementation costs*. As pointed out in the previous paragraph, implementing a new specification for exchanging patient medical record information will require significant effort regardless of whether version 2.x or version 3 is considered. We believe that using version 3 will have a lower cost in the long term, as the specificity of the standard itself will enable implementers to reuse development efforts with greater ease than is currently possible with version 2.x. Based on our experience, we also believe that properly engineered applications will have little difficulty switching from version 2.x to version 3 over a period of several years. Indirectly, adopting version 3 transaction sets as a standard specification will improve competition in the market place, as better engineered applications will have an appropriate advantage.

In conclusion, we would like to thank NCVHS for giving us the opportunity to present our work and ideas on an industry standard for the PMRI structure. We hope that such a standard will be finalized in the near future and that vendors, like us, contribute to the success of this effort with a fast and thorough implementation of the standard.